Performance of Some CFD Codes on the Alliant FX/8

Walid Abu-Sufah
Amjad Daoud

TR 88-11

PERFORMANCE OF SOME * CFD CODES ON THE ALLIANT FX/8

Presented at the third SIAM conference
on Parallel Processing dor Scientific Computing
Dec 1-4,1987

Walid Abu-Sufah
Amjad Daoud

Department of Computer Science
Virginia Polytechnic Institute and State University
Blacksburg , VA24061

* The empirical data was collected using the Advanced Computing Research Faciltiy of the Argonne National Lab.

Environment

Work performed on the Alliant FX/8
at the ACRF of Argonne National Lab.

Configuration :

8	Computational Processors (CE's)
6	Interactive Processors
128	Kbytes Shared Cache
32	Mbytes Shared Memory

CFD Codes

1) The NAS Kernel Benchmark Program

2) AERO40 : 3-D Unsteady Euler Equations solver using an explicit finite-volume Runge-Kutta time stepping method.

3) AERO60 : Solves the 3-D Unsteady Euler equations using the Beam and Worming Implicit finite difference method.

4) ARC2D : 2D Navier-Stokes Equations using implicit method.

Kernels Description

1. **MXM** - This subroutine performs the usual matrix product on two input matrices. The subroutine employs a four-way unrolled, outer product matrix multiply algorithm that is especially effective for most vector computers.

2. **CFFT2D** - This test performs a complex radix 2 FFT on a two dimensional input array, returning the result in place. The test kernel actually consists of two subroutines that perform FFTs along the first and second dimension of the array, respectively, taking advantage of the parallel structure of the array.

3. **CHOLSKY** - This subroutine performs a Cholesky decomposition in parallel on a set of input matrices, which are actually input to the subroutine as a single three dimensional array.

4- **BTRIX** - This kernel performs a block tridiagonal matrix solution along one dimension of a four dimensional array.

5- **GMTRY**- This subroutine sets up arrays for a vortex method solution and performs Gaussian elimination on the resulting array. This kernel is noted for a number of loops that are challenging to vectorize.

6- **EMIT** - Also extracted from a vortex code, this subroutine creates new vortices according to certain boundary conditions.

7- **VPENTA** - This subroutine simultaneously inverts three matrix pentadiagonals in a highly parallel fashion.

Table 1: Kernel Features

Feature	KERNEL						
	1	2	3	4	5	6	7
Two dimensional arrays	x	x			x	x	x
Multidimensional arrays			x	x			x
Dimensions with colons			x				
Integer arrays		x			x	x	
Integer functions in indices					x	x	
IF statements in inner loops						x	
Scientific function calls		x	x		x	x	
Complex arithmetic		x			x	x	
Complex function calls					x	x	
Inner loop memory strides	1	1 2 256	1 4	1 2 750 900	1 2 500	1	128
Inner loop vector lengths	256	128 256	250	28	5 100 500	100 500 1000	128

The Alliant FX/8

Multiprocessor Performance

(Level Zero Tuning)

TABLE 2 *

NAS KERNEL BENCHMARK RESULTS (MFLOPS)

Computer System	No. CPU	Tuning	Kernels 1	2	3	4	5	6	7	Comp Rate
Cray X-MP /48	1	0	136.0	45.9	59.8	82.3	95.5	84.1	30.5	61.9
Cray X-MP /48	1	20	136.0	85.2	66.7	79.6	115.5	103.0	124.1	96.4
Cray X-MP /48	4	20	536.8	330.9	205.0	273.3	395.3	396.6	483.9	349.1
Cray-2	1	0	142.4	9.0	6.1	42.7	4.0	75.0	5.9	10.6
Cray-2	1	20	193.0	58.1	48.2	66.0	74.4	141.4	135.6	80.9
Cray-2	1	50	193.2	137.1	48.5	65.6	75.2	141.0	135.9	99.0
Cray-2	4	50	590.0	317.2	110.2	162.1	233.1	544.7	312.7	253.8
CDC 205	1	0	116.6	12.5	24.2	8.0	21.3	61.1	9.4	16.1
CDC 205	1	20	129.8	49.5	108.4	14.5	72.1	76.9	52.8	44.7
CDC 205	1	50	127.8	57.4	108.3	135.7	75.0	76.2	67.4	82.9
Amdahl 1200	1	0	465.1	11.1	42.2	88.5	38.3	214.5	7.3	22.4
Amdahl 1200	1	20	497.2	106.0	95.6	88.0	127.5	214.9	202.3	139.1
Amdahl 1200	1	50	500.9	106.5	96.1	91.3	127.4	220.5	202.4	140.8
Amdahl 1200	1	unl	499.2	162.1	96.7	124.5	150.9	219.4	232.2	174.7

TABLE 2 * (CONTINUED)

NAS KERNEL BENCHMARK RESULTS (MFLOPS)

Computer System	No. CPU	Tuning	Kernels 1	2	3	4	5	6	7	Comp Rate
Amdahl 1400	1	0	763.9	12.2	71.7	101.3	61.6	404.1	15.6	31.6
Amdahl 1400	1	20	762.6	166.3	118.7	100.6	146.7	402.9	221.9	182.8
Amdahl 1400	1	50	698.3	169.0	118.6	101.9	146.9	402.4	221.8	183.4
Amdahl 1400	1	unl	692.0	203.7	119.5	155.1	147.1	418.8	221.8	212.4
NEC SX-2	1	0	687.0	20.5	281.3	147.0	273.4	301.1	22.4	52.9
NEC SX-2	1	20	821.8	308.7	319.2	147.0	277.0	437.6	384.9	309.7
FPS-264	1	0	16.3	8.5	3.5	7.6	4.3	6.7	8.0	7.0
CONVEX C-1 XP	1	20	15.8	5.7	4.0	4.0	6.9	6.4	6.6	6.1
ALLIANT FX/8	1	0	2.94	1.20	1.27	1.60	0.53	4.06	0.85	1.26
ALLIANT FX/8	4	0	10.92	1.91	3.20	3.82	1.68	7.94	1.28	2.65
ALLIANT FX/8	8	0	18.48	1.89	3.66	5.56	3.12	13.72	1.71	3.37

* All but the Alliant FX/8 numbers are taken from NASA/Ames (June 2, 1987)

TABLE 3

NAS KERNEL BENCHMARK RESULTS

(Execution rate / Peak execution rate)

Computer System	No. CPU	Tun-ing	Kernels 1	2	3	4	5	6	7	Comp Rate
Cray X-MP /48	1	0	0.65	0.22	0.28	0.39	0.46	0.40	0.14	0.30
Cray X-MP /48	1	20	0.65	0.41	0.32	0.38	0.55	0.49	0.59	0.46
Cray X-MP /48	4	20	0.64	0.39	0.24	0.33	0.47	0,47	0.58	0.42
Cray-2	1	0	0.29	0.02	0.01	0.09	0.008	0.15	0.01	0.02
Cray-2	1	20	0.40	0.12	0.10	0.14	0.15	0.29	0.28	0.17
Cray-2	1	50	0.40	0.28	0.10	0.13	0.15	0.29	0.28	0.20
Cray-2	4	50	0.30	0.16	0.06	0.08	0.12	0.28	0.16	0.13
CDC 205	1	0	0.29	0.03	0.06	0.02	0.05	0.15	0.02	0.04
CDC 205	1	20	0.32	0.12	0.27	0.04	0.18	0.19	0.13	0.11
CDC 205	1	50	0.32	0.14	0.27	0.34	0.19	0.19	0.17	0.21
Amdahl 1200	1	0	0.87	0.02	0.08	0.17	0.07	0.40	0.01	0.04
Amdahl 1200	1	20	0.93	0.20	0.18	0.17	0.24	0.40	0.38	0.26
Amdahl 1200	1	50	0.94	0.20	0.18	0.17	0.24	0.41	0.38	0.26
Amdahl 1200	1	unl	0.94	0.30	0.18	0.23	0.28	0.41	0.44	0.33

TABLE 3 (Continued)

NAS KERNEL BENCHMARK RESULTS

(Execution rate/ Peak Execution rate)

Computer System	No. CPU	Tun-ing	Kernels							Comp Rate
			1	2	3	4	5	6	7	
Amdahl 1400	1	0	0.65	0.01	0.06	0.09	0.05	0.35	0.01	0.03
Amdahl 1400	1	20	0.67	0.15	0.10	0.09	0.13	0.35	0.19	0.16
Amdahl 1400	1	50	0.61	0.15	0.10	0.09	0.13	0.35	0.19	0.16
Amdahl 1400	1	unl	0.61	0.18	0.11	0.14	0.13	0.37	0.19	0.19
NEC SX-2	1	0	0.53	0.02	0.22	0.11	0.21	0.23	0.02	0.04
NEC SX-2	1	20	0.63	0.24	0.25	0.11	0.21	0.34	0.30	0.24
FPS-264	1	0	0.43	0.22	0.09	0.20	0.11	0.18	0.21	0.18
CONVEX C-1 XP	1	20	0.79	0.29	0.20	0.20	0.35	0.32	0.33	0.31
ALLIANT FX/8	1	0	0.50	0.20	0.22	0.27	0.09	0.69	0.14	0.21
ALLIANT FX/8	4	0	0.46	0.08	0.14	0.16	0.07	0.34	0.05	0.11
ALLIANT FX/8	8	0	0.39	0.04	0.08	0.12	0.07	0.29	0.04	0.07

NAS RESULTS

SPEEDUP RESULTS

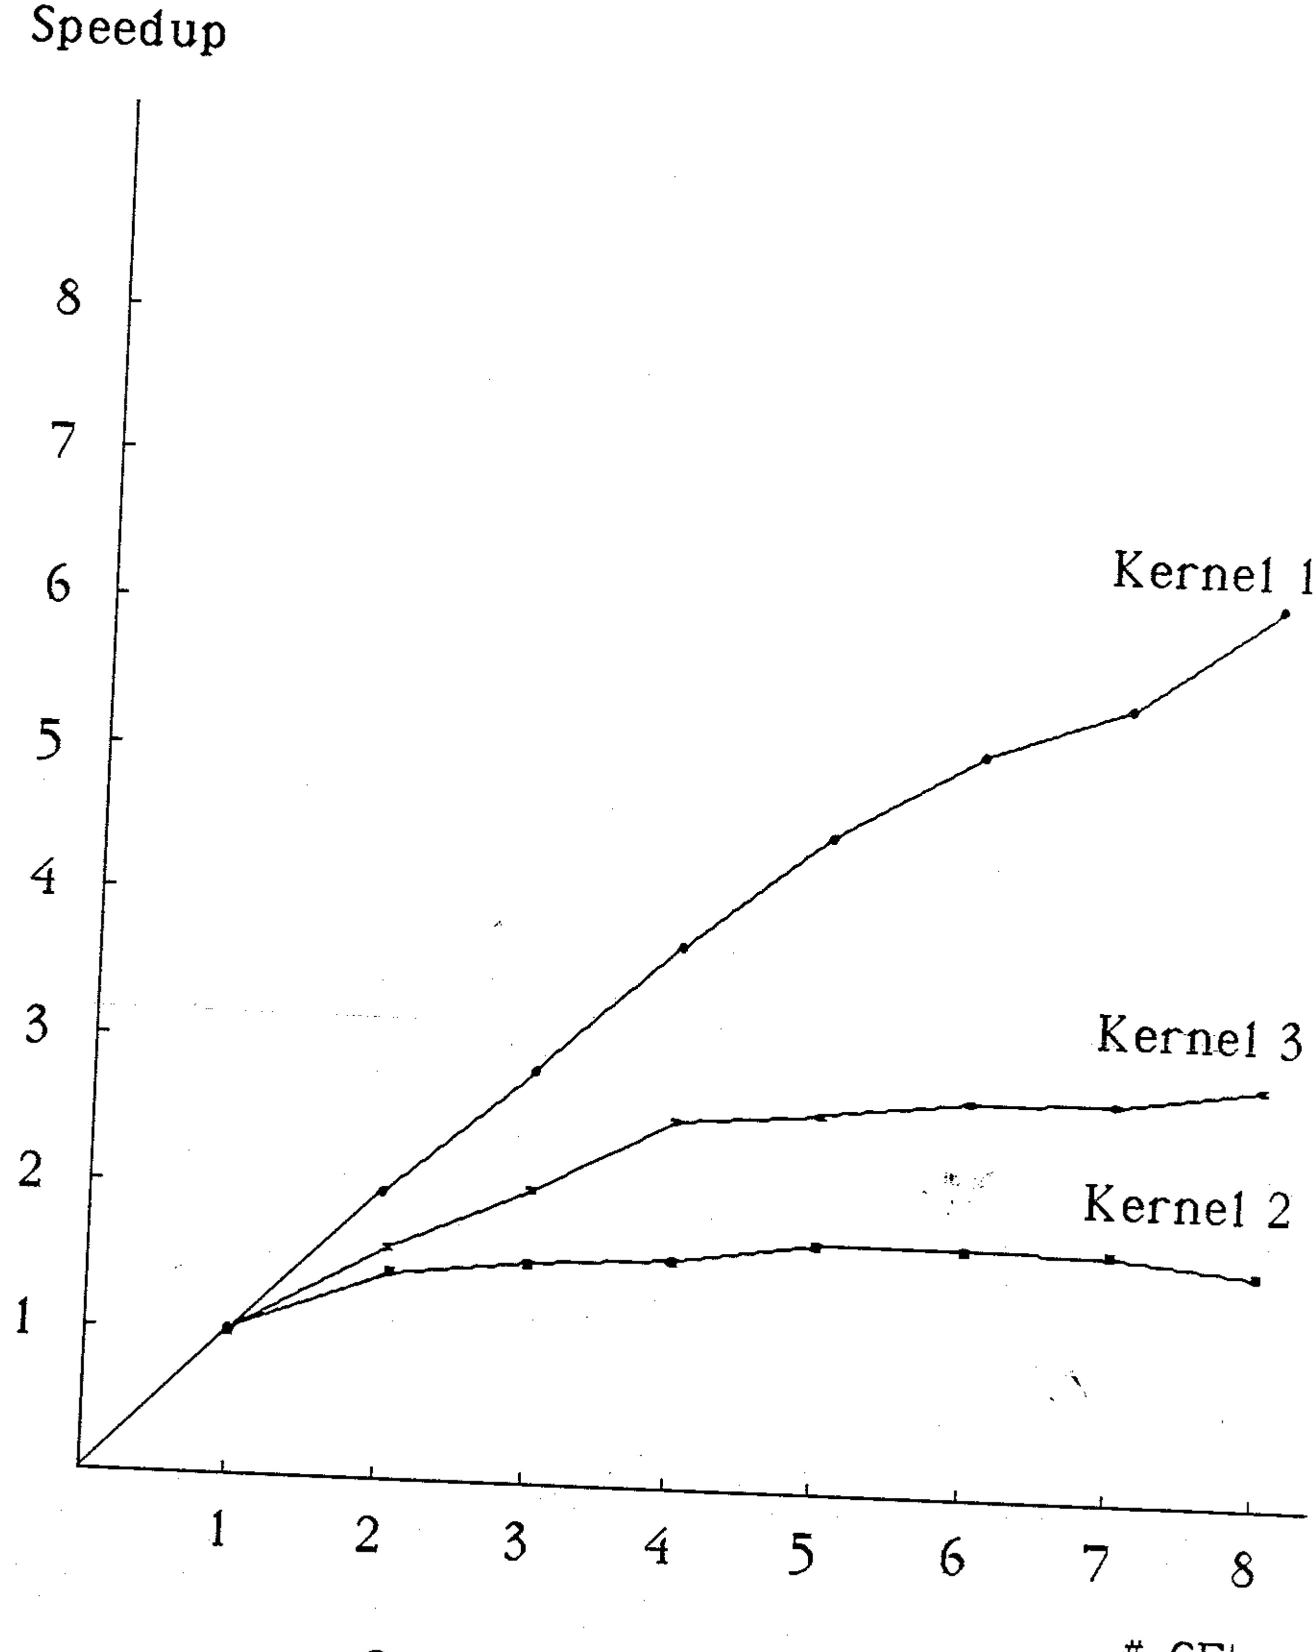

Speedup vs. #CE's

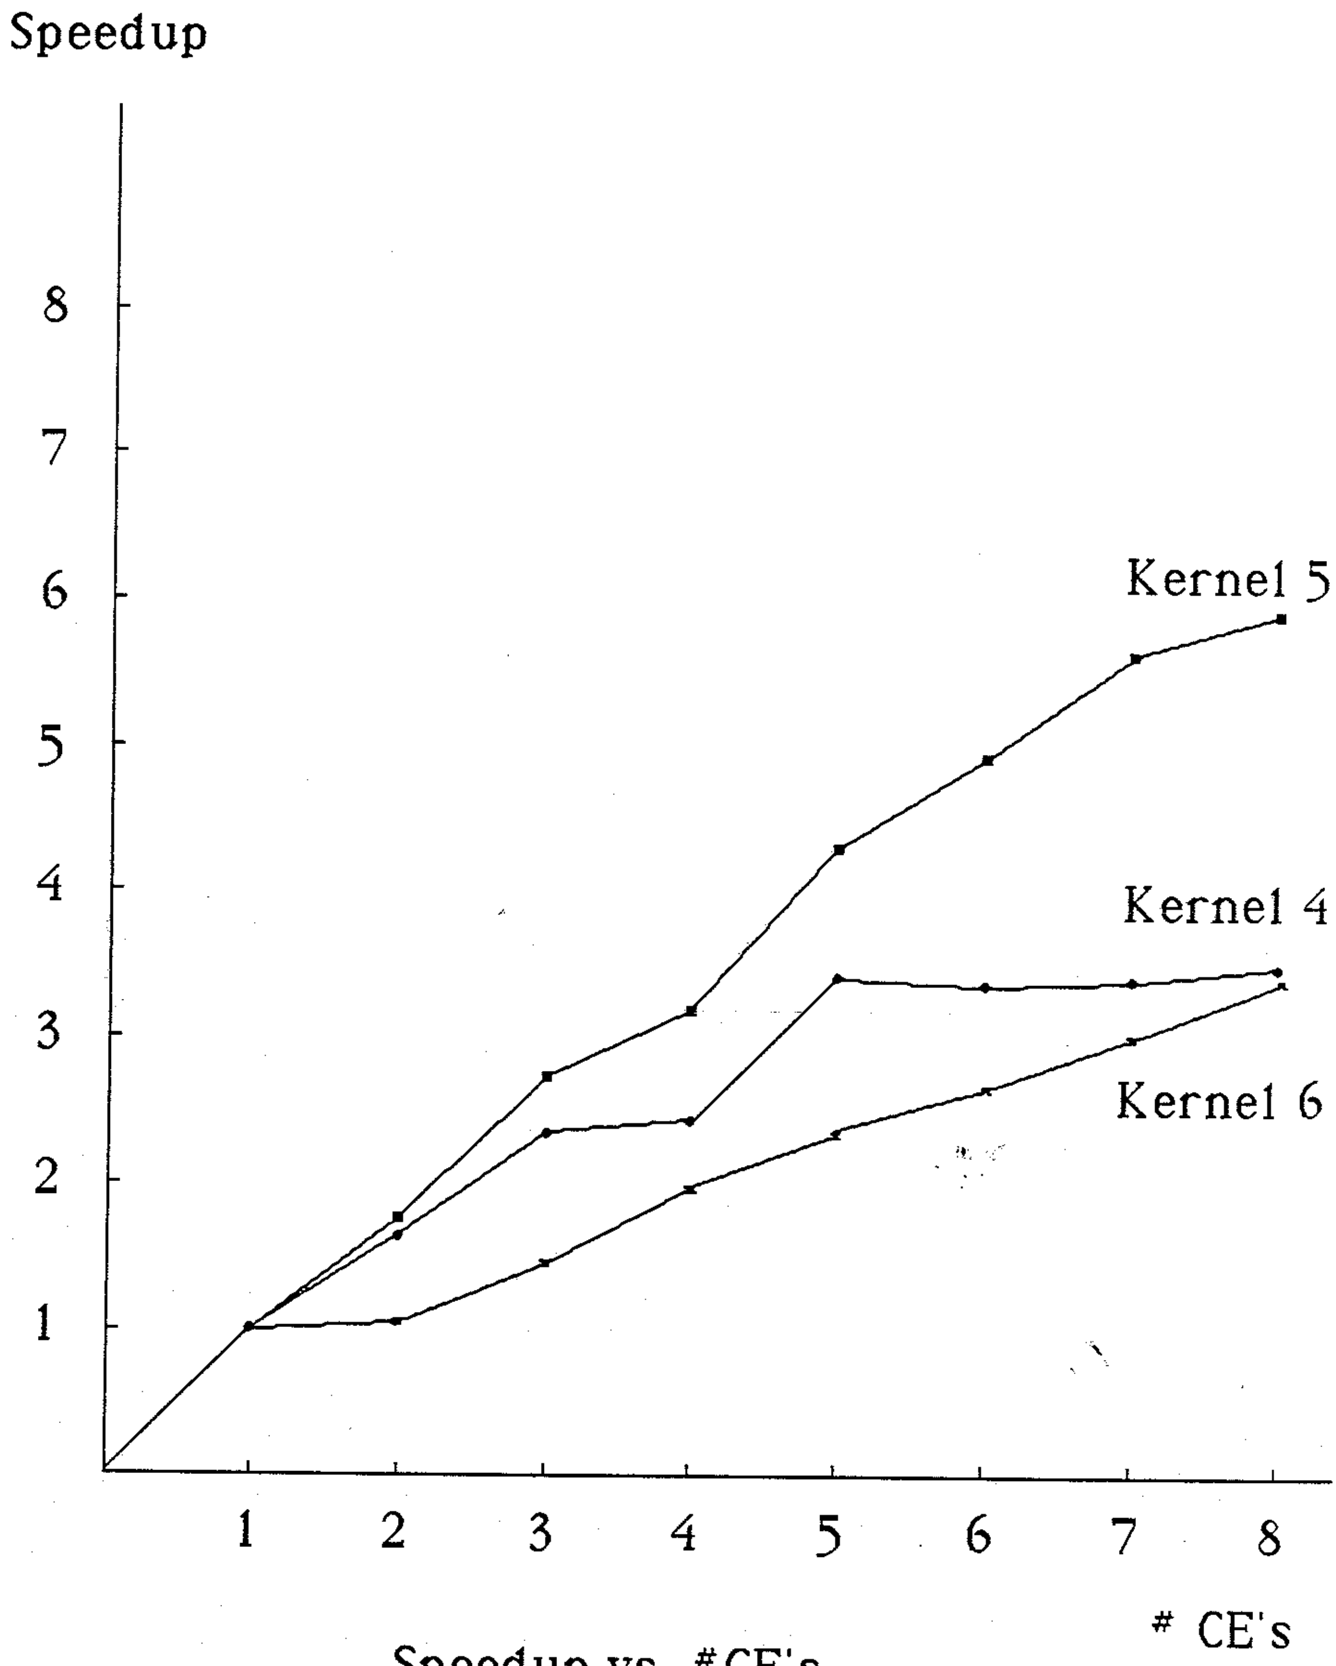

Speedup vs. #CE's

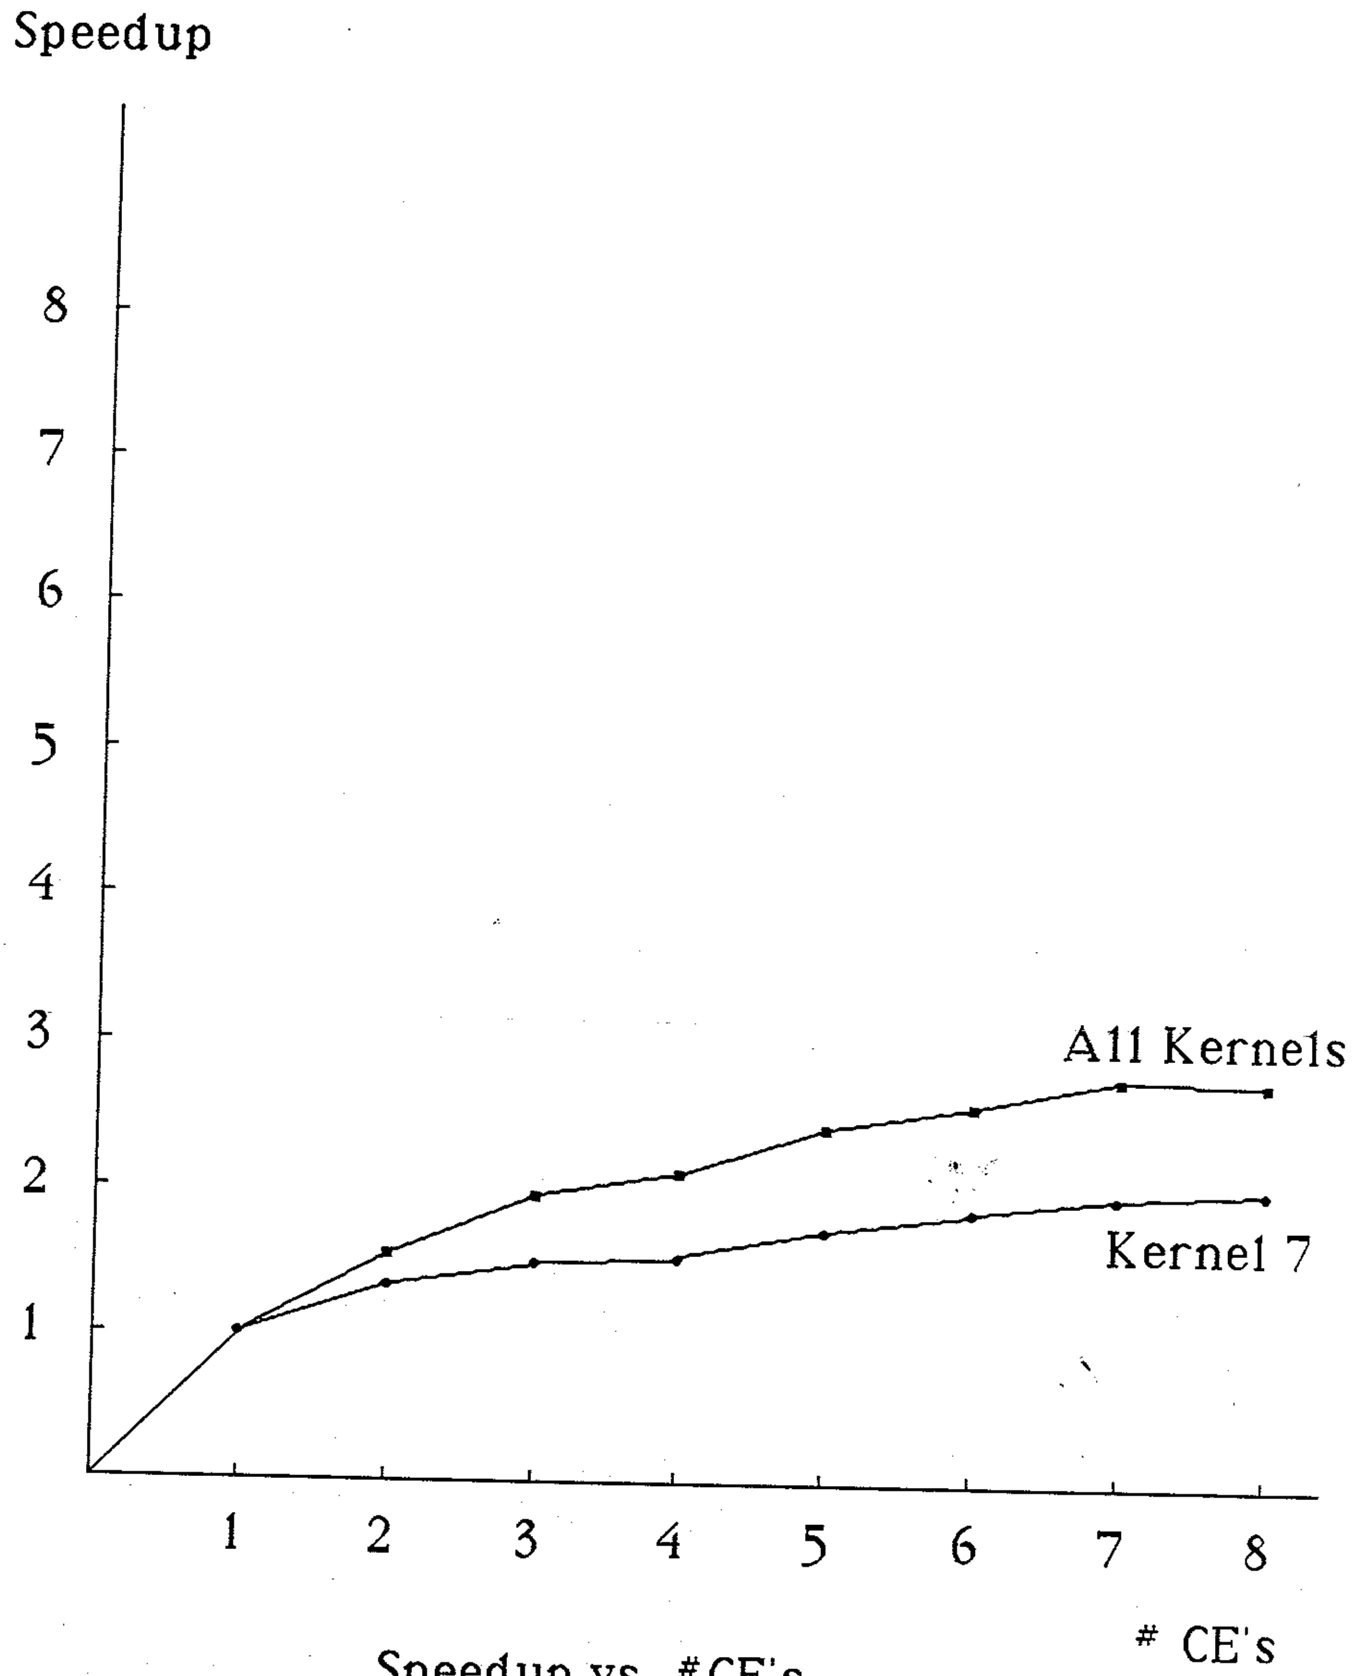

Speedup vs. #CE's

NAS Kernels

Multiprocessing Efficiency

Results

$$E_p = \text{Speedup} / \# \text{ of CE's}$$

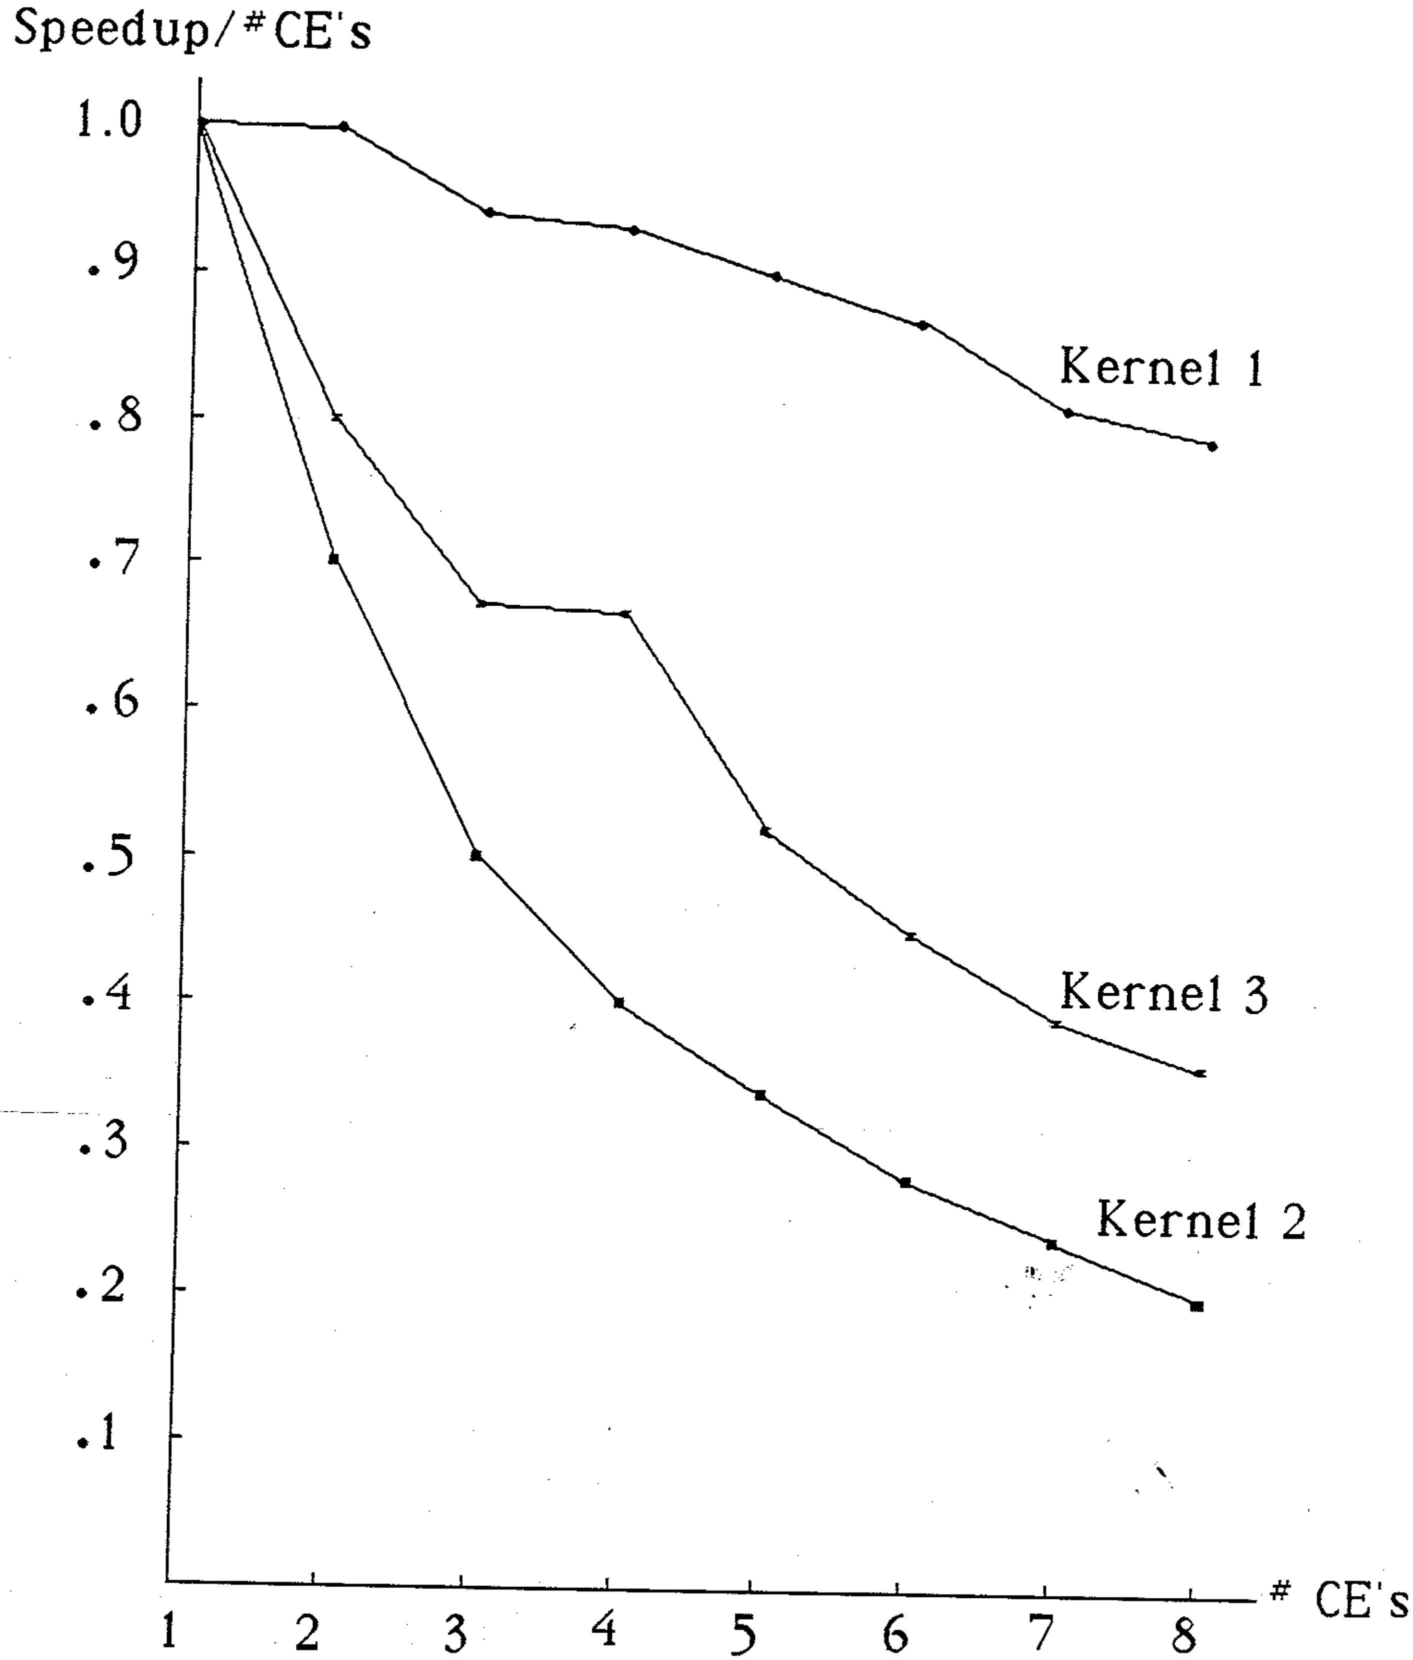

Multiprocessing Efficiency vs. # of CE's

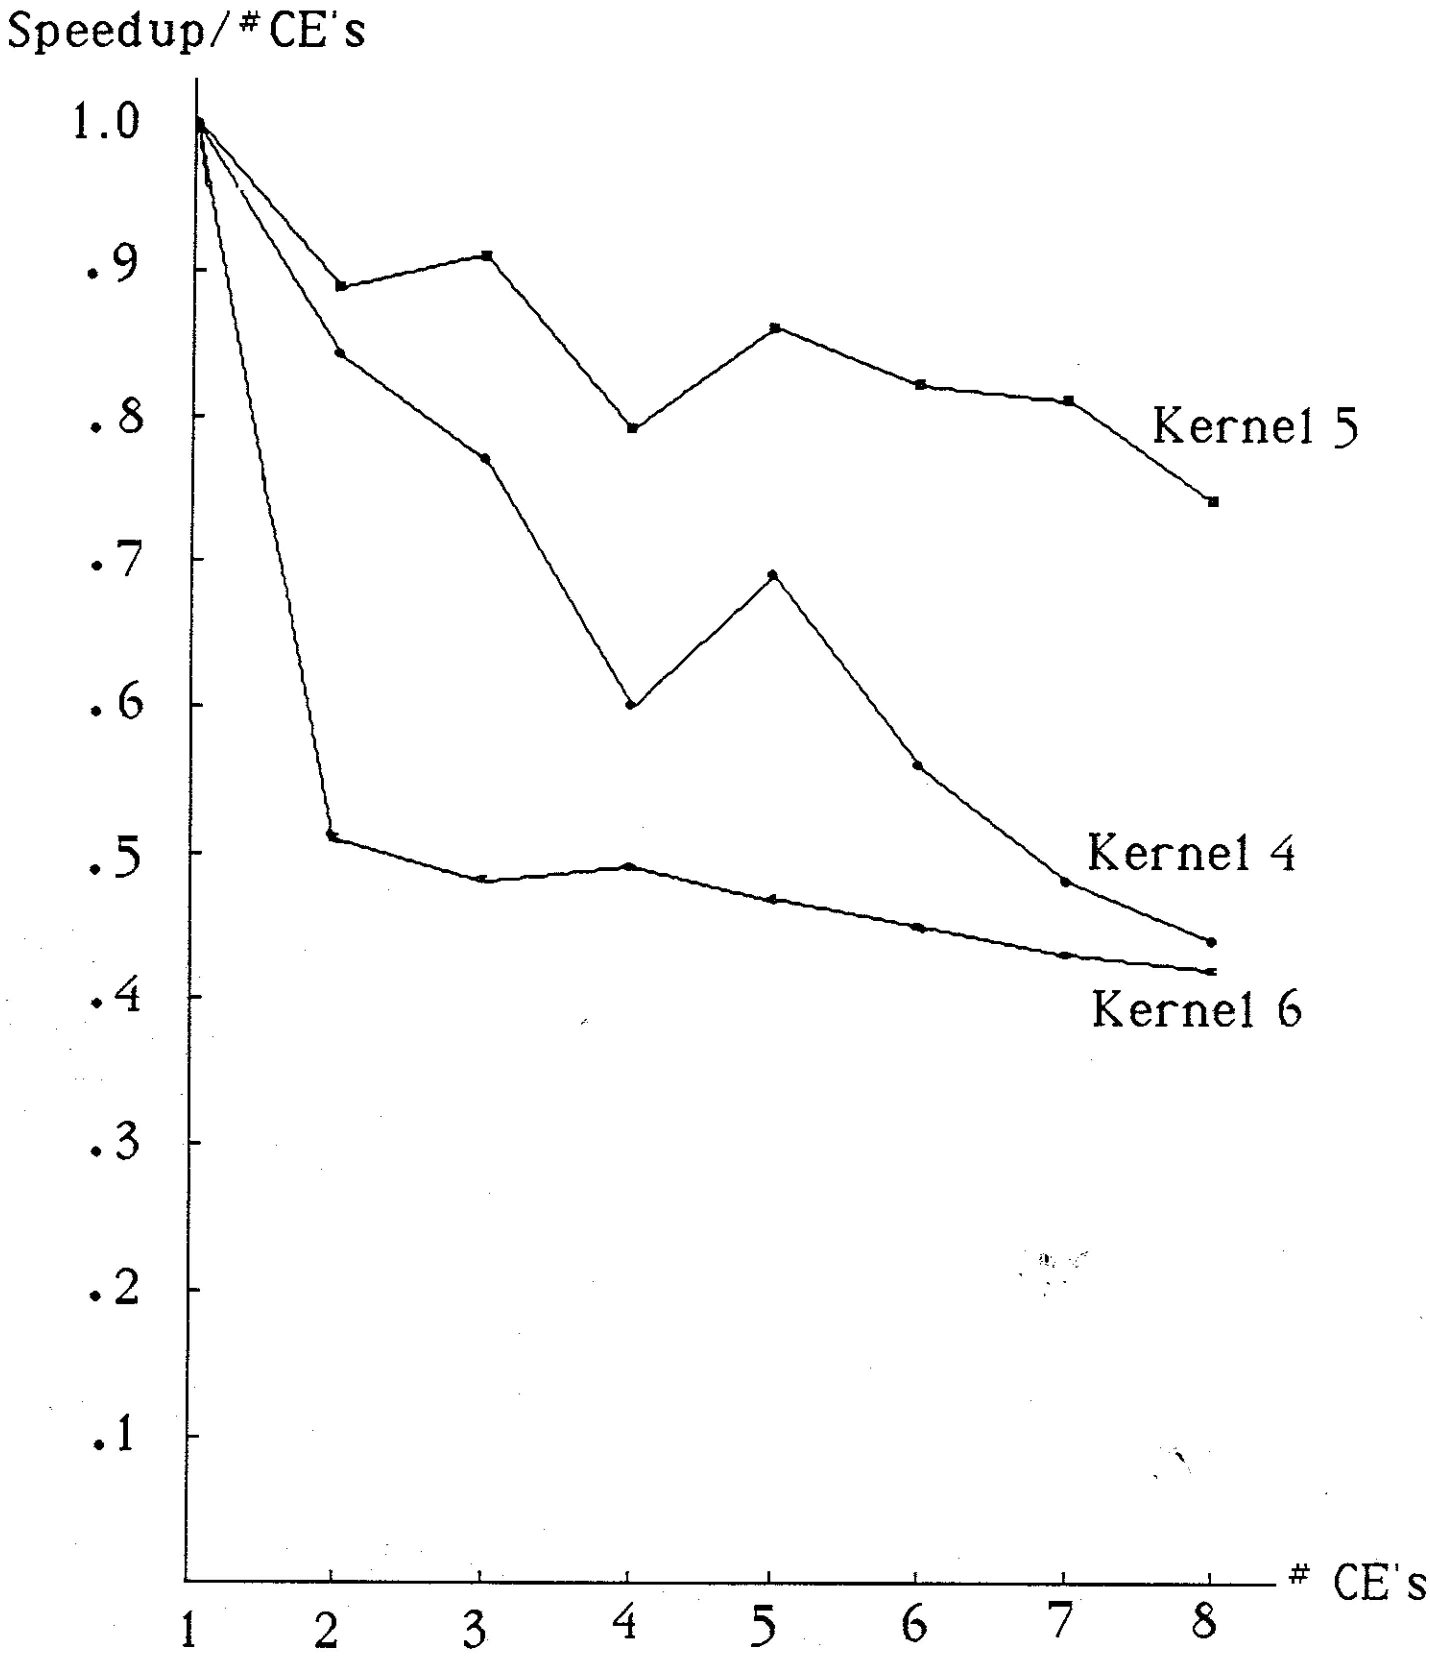

Multiprocessing Efficiency vs. # of CE's

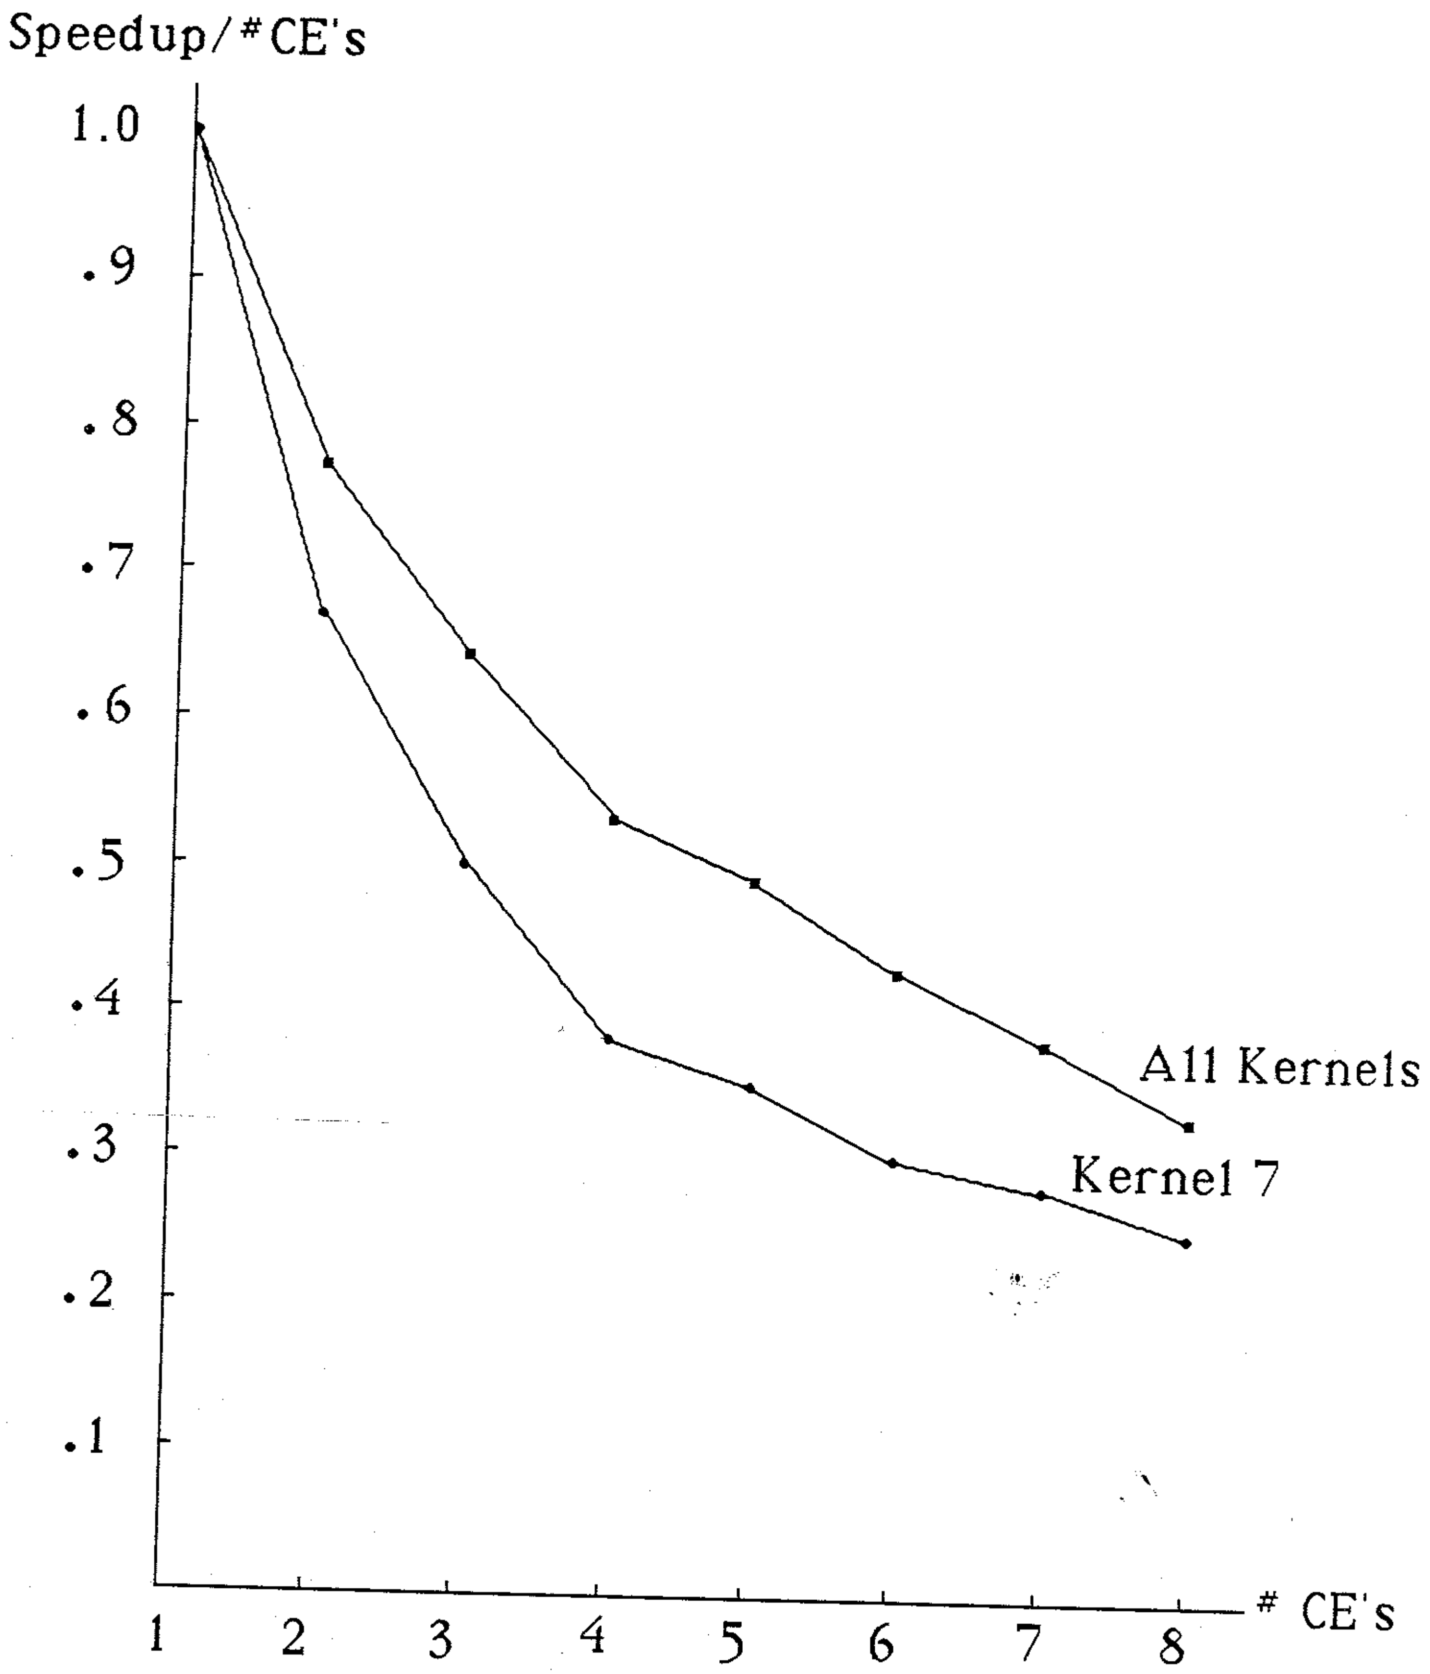

Multiprocessing Efficiency vs. # of CE's

NAS Kernels Results
on
Alliant FX/8
(Zero Level Tuning)

1. A single CE seems to be a well designed vector processor.

2. The vectorizing compiler is good

3. Multiprocessing performance levels off at 4 CE's.

4. Computational processors bandwidth versus memory system bandwidth !!

5. Vector length is divided on Several CE's => higher pipelining overhead.

Performance
Metrics Questions

- MFLOPS ?!!
- Machine Efficiency ?!!
- Speedup ?!!
- Multiprocessing Efficiency !!??
  { Academic }

AERO40 RESULTS

{ Work still in <u>PROGRESS</u> }

Aero40 Results
Speedup

A. t_0 = execution time of code with NO optimization

t_{18} = execution time with -O on 8 CE's

Speedup = $S_{18} = 26.4$

B. t_2 = execution with -g optimization

Speedup = $t_0/t_2 = 4.93$

C. t_3 = execution with -Ogv optimization on a single CE

Speedup (1 CE vector) = $t_0/t_3 = 9.04$

D. t_4 = execution with -Ogc optimization
-> concurrent execution on 8 CE's

Speedup = $t_0/t_4 = 24.4$

{ with vectorization 26.4 }

More Speedup Results

E. $t_5 =$ time on one CPU (CE) with vectorization

$t_6 =$ 8 CE's execution time, concurrency optimization

Speedup $= t_5/t_6 = 5.20$

F. $t_7 =$ 8 CE's vectorization and concurrency

Speedup $= t_5/t_7 = 3.01$

EULER CODE

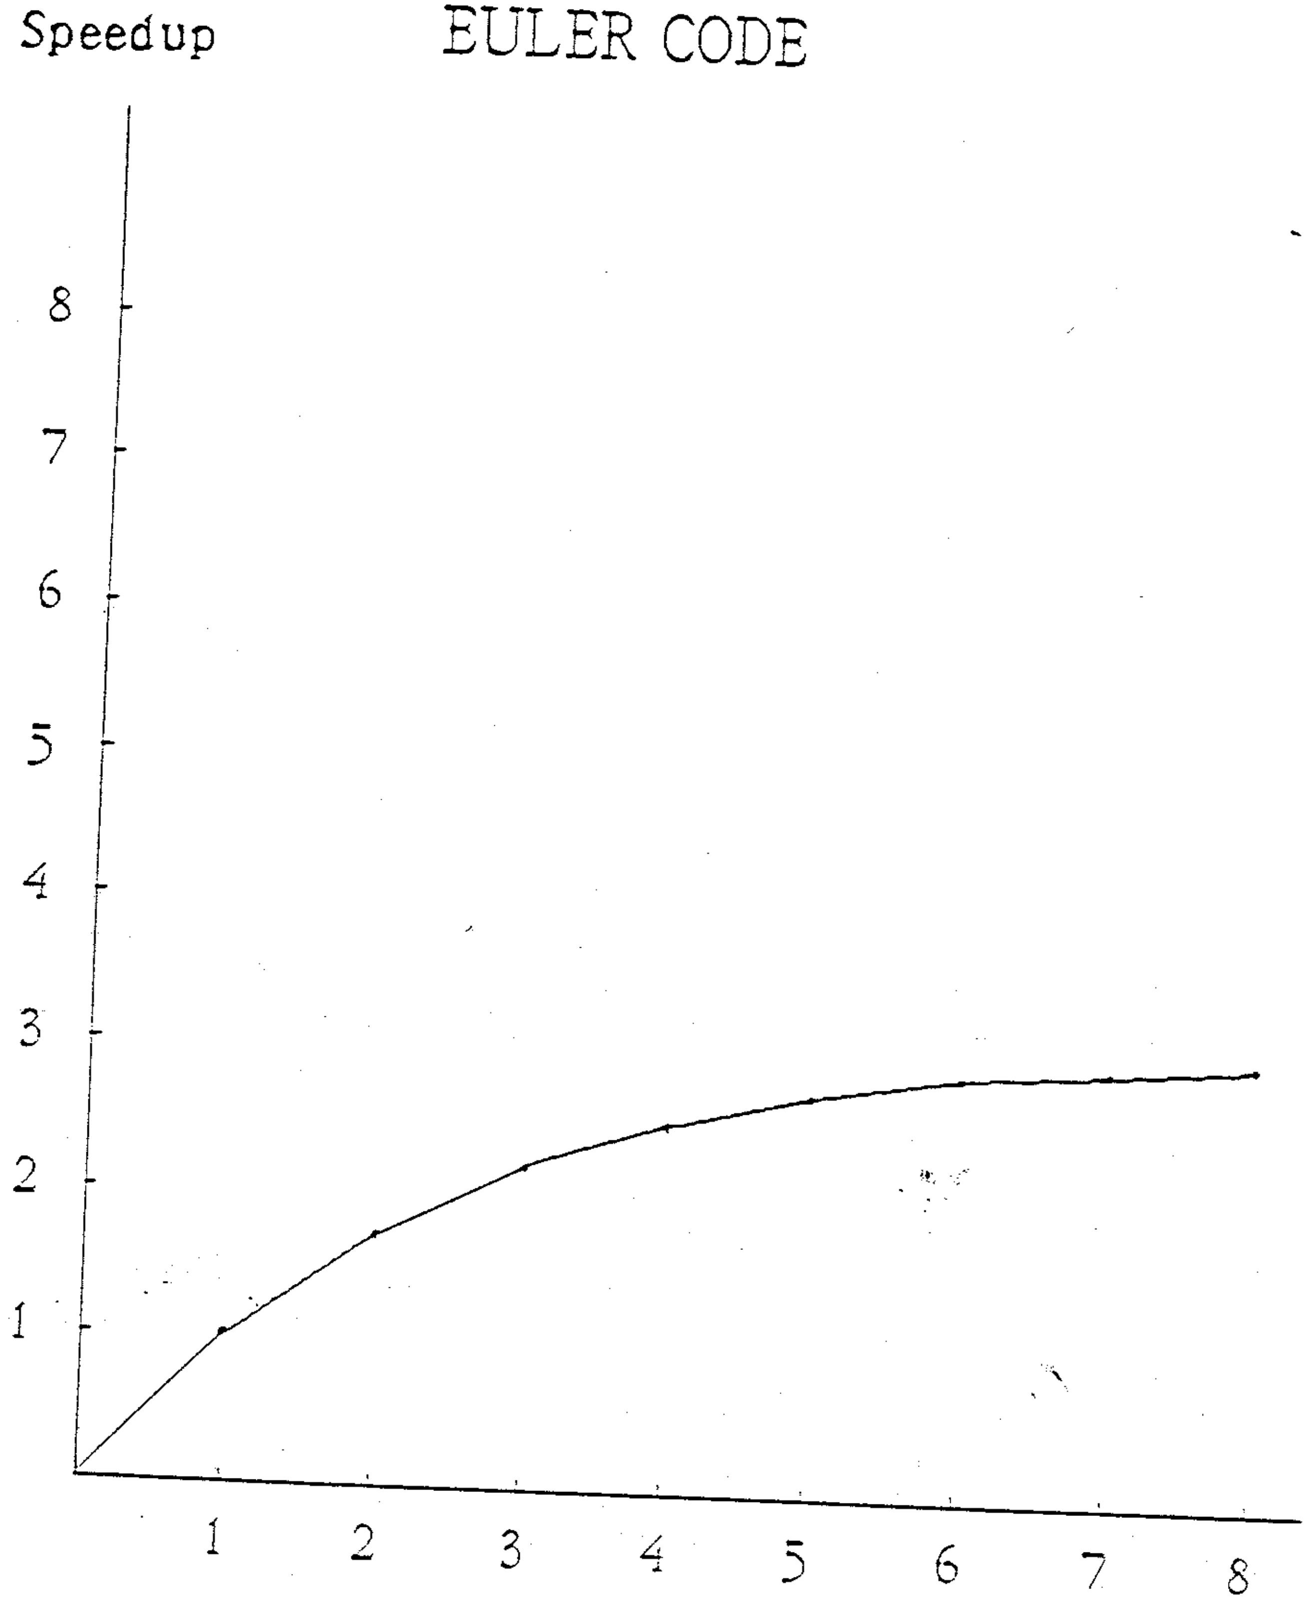

Speedup vs. # CE's

EULER CODE

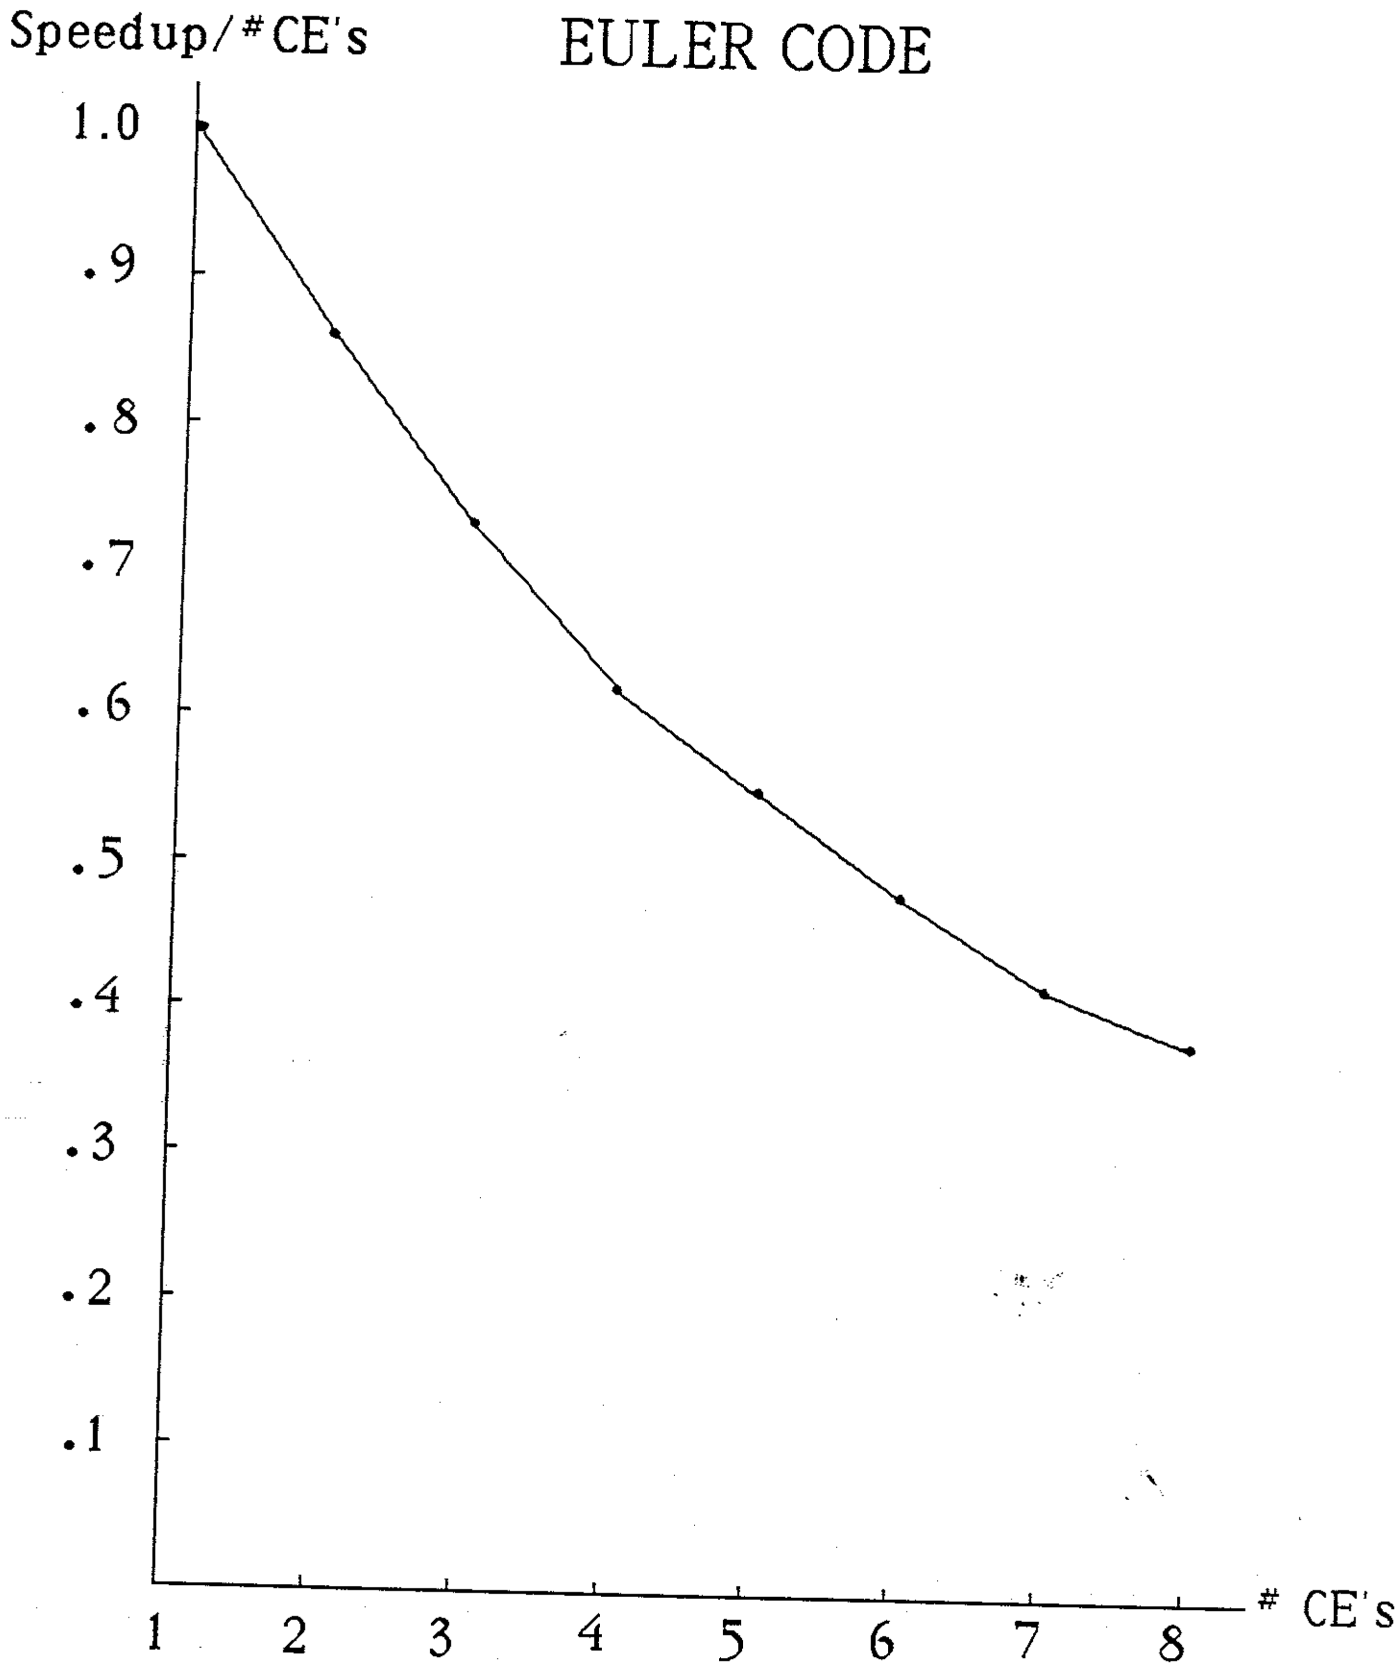

Multiprocessing Efficiency vs. # of CE's

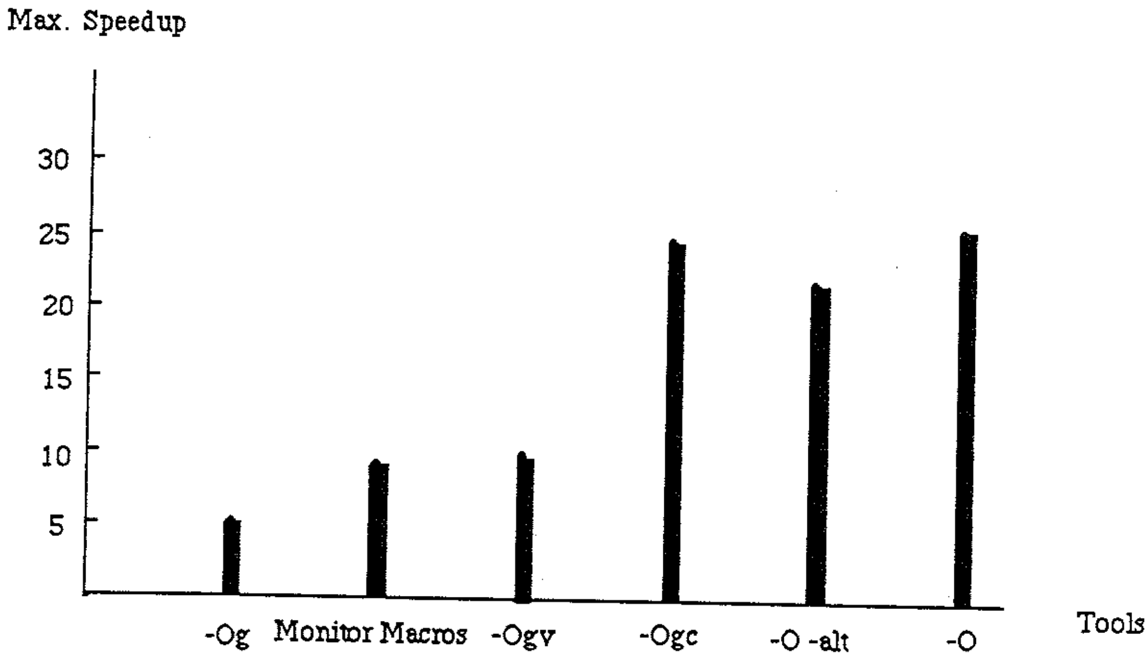

Figure 6

Maximum Speedup obtained using different tools or options

Initial Conclusions

AERO40 Results

*Memory hierarchy considerations are the basic reasons for the obtained speedup while using parallelism and vectorization on 8 CE's

* Not enough memory system bandwidth for 8 CE's

* Needed :

- Memory hierarchy management transformations in the preprocessor.

- Memory hierarchy management directives to enable the user to advise the compiler on better memory management.